

WHAT IS CLAIMED IS:

1 1. A system for optimizing data access, comprising:
2 a file server capable of communicating with one or more clients; and
3 a plurality of storage elements organized into pairs for storing a plurality of
4 files, each pair having a master storage element and one or more mirrored storage elements,
5 and each mirrored storage element having a copy of data stored on the master storage
6 elements

7 wherein the file server maintains file information on where each of the
8 plurality of files is stored on which pair of storage elements;

9 wherein the file server further maintains access load information on each pair
0 of storage elements; and

1 wherein when a client requests file information for a requested file from the
2 file server, the file server determines which pair of storage elements has the requested file and
3 further determines which of the storage elements within the pair of storage elements having
4 the requested file is to be accessed.

1 2. The system according to claim 1 wherein the plurality of storage
2 elements is a plurality of disk drives.

1 3. The system according to claim 1 wherein the plurality of storage
2 elements are stored on a single storage system.

1 4. The system according to claim 1 wherein the plurality of storage
2 elements are stored on one or more storage systems.

1 5. The system according to claim 1 further comprising:
2 a plurality of host computers;
3 wherein the file server resides on one of the plurality of host computers; and
4 wherein the one or more clients reside on remaining ones of the plurality of
5 host computers.

1 6. The system according to claim 1 further comprising:
2 a sync daemon configured to synchronize data stored on each pair of storage
3 elements.

1 7. The system according to claim 1 wherein the master storage element
2 and the one or more mirrored storage elements within a pair are stored on a single storage
3 system.

1 8. The system according to claim 1 wherein the master storage element
2 and the one or more mirrored storage elements within a pair are stored on one or more
3 storage systems.

1 9. The system according to claim 8 wherein if it is determined that a
2 mirrored storage element is to be accessed for the requested file and the mirrored storage
3 element which is to be accessed contains a latest copy of data for the requested file stored on
4 the corresponding master storage element, the client directly retrieves the requested file from
5 the mirrored storage element.

1 10. The system according to claim 8 wherein if it is determined that a
2 mirrored storage element is to be accessed for the requested file and the mirrored storage
3 element which is to be accessed does not contain a latest copy of data for the requested file
4 stored on the corresponding master storage element, the latest copy of data for the requested
5 file stored on the corresponding master storage element is retrieved from the corresponding
6 master storage element and then forwarded to the client.

1 11. The system according to claim 1 wherein the file information on where
2 each of the plurality of files is stored on which pair of storage elements includes file
3 allocation lists.

1 12. The system according to claim 1 wherein when determining which of
2 the storage elements within the pair of storage elements having the requested file is to be
3 accessed, consideration is given to ensure that all the storage elements within the pair of
4 storage elements having the requested file are substantially accessed in a balanced manner.

1 13. The system according to claim 1 wherein upon determining which of
2 the storage elements within the pair of storage elements having the requested file is to be
3 accessed, the file server forwards information relating to the determination to the client
4 thereby allowing the client to retrieve the requested file from the determined storage element.

1 14. The system according to claim 13 wherein upon forwarding the
2 information relating to the determination to the client, the file server updates the access load
3 information to ensure accurate monitoring of access balance of the pairs.

1 15. A system for optimizing data access, comprising:
2 a first host computer having a file system server, the first host computer
3 capable of communicating with a second host computer having a file system client; and
4 a storage system having a plurality of disk drives organized into pairs for
5 storing a plurality of files, each pair having a master disk drive and one or more mirrored disk
6 drives, each mirrored disk drive having a copy of data stored on the master disk drive;
7 wherein the file system server maintains file information on where each of the
8 plurality of files is stored on which pair of disk drives;
9 wherein the file system server further maintains access load information on
10 each pair of disk drives; and
11 wherein when a file system client requests file information for a requested file
12 from the file system server, the file system server determines which pair of disk drives has
13 the requested file and further determines which of the disk drives within the pair of disk
14 drives having the requested file is to be accessed.

1 16. The system according to claim 15 wherein the file information on
2 where each of the plurality of files is stored on which pair of disk drives includes file
3 allocation lists.

1 17. The system according to claim 15 wherein when determining which of
2 the disk drives within the pair of disk drives having the requested file is to be accessed,
3 consideration is given to ensure that all the disk drives within the pair of disk drives having
4 the requested file are substantially accessed in a balanced manner.

1 18. The system according to claim 15 wherein upon determining which of
2 the disk drives within the pair of disk drives having the requested file is to be accessed, the
3 file system server forwards information relating to the determination to the file system client
4 thereby allowing the file system client to retrieve the requested file from the determined disk
5 drive.

19202122232425262728292A
TIECXO 441660

1 19. The system according to claim 18 wherein upon forwarding the
2 information relating to the determination to the file system client, the file system server
3 updates the access load information to ensure accurate monitoring of access balance of the
4 pairs.

1 20. A system for optimizing data access, comprising:
2 a first host computer having a file system server, the first host computer
3 capable of communicating with a second host computer having a file system client; and
4 a plurality of storage systems, each of the plurality of storage systems having a
5 plurality of disk drives, the plurality of disk drives from the plurality of storage systems are
6 collectively organized into pairs for storing a plurality of files, each pair having a master disk
7 drive and one or more mirrored disk drives, each mirrored disk drive having a copy of data
8 stored on the master disk drive;
9 wherein the file system server maintains file information on where each of the
10 plurality of files is stored on which pair of disk drives;
11 wherein the file system server further maintains access load information on
12 each pair of disk drives; and
13 wherein when a file system client requests file information for a requested file
14 from the file system server, the file system server determines which pair of disk drives has
15 the requested file and further determines which of the disk drives within the pair of disk
16 drives having the requested file is to be accessed.

1 21. The system according to claim 20 wherein the file information on
2 where each of the plurality of files is stored on which pair of disk drives includes file
3 allocation lists.

1 22. The system according to claim 20 wherein when determining which of
2 the disk drives within the pair of disk drives having the requested file is to be accessed,
3 consideration is given to ensure that all the disk drives within the pair of disk drives having
4 the requested file are substantially accessed in a balanced manner.

1 23. The system according to claim 20 wherein upon determining which of
2 the disk drives within the pair of disk drives having the requested file is to be accessed, the
3 file system server forwards information relating to the determination to the file system client

4 thereby allowing the file system client to retrieve the requested file from the determined disk
5 drive.

1 24. The system according to claim 20 wherein upon forwarding the
2 information relating to the determination to the file system client, the file system server
3 updates the access load information to ensure accurate monitoring of access balance of the
4 pairs.

1 25. The system according to claim 20 wherein if it is determined that a
2 mirrored disk drive is to be accessed for the requested file and the mirrored disk drive which
3 is to be accessed contains a latest copy of data for the requested file stored on the
4 corresponding master disk drive, the file system client directly retrieves the requested file
5 from the mirrored disk drive.

1 26. The system according to claim 20 wherein if it is determined that a
2 mirrored disk drive is to be accessed for the requested file and the mirrored disk drive which
3 is to be accessed does not contain a latest copy of data for the requested file stored on the
4 corresponding master disk drive, the latest copy of data for the requested file stored on the
5 corresponding master disk drive is retrieved from the corresponding master disk drive and
6 then forwarded to the file system client.

1 27. A method for optimizing data access, comprising:
2 organizing a plurality of storage elements into pairs for storing a plurality of
3 files, each pair having a master storage element and one or more mirrored storage elements,
4 and each mirrored storage element having a copy of data stored on the master storage
5 element;
6 maintaining file information on where each of the plurality of files is stored on
7 which pair of storage elements;
8 maintaining access load information on each pair of storage elements; and
9 upon receiving a request for a requested file, determining which pair of
10 storage elements has the requested file by using the file information, and then determining
11 which of the storage elements within the pair of storage elements having the requested file is
12 to be accessed by using the access load information.

1 28. The method of claim 27 further comprising:

2 upon determining which of the storage elements within the pair of storage
3 elements having the requested file is to be accessed, forwarding information relating to the
4 determination to a client thereby allowing the client to retrieve the requested file.

1 29. The method of claim 28 further comprising:
2 upon forwarding the information relating to the determination to the client,
3 updating the access load information to ensure accurate monitoring of access balance of the
4 pairs.

1 30. The method according to claim 27 wherein the plurality of storage
2 elements are stored on a single storage system.

1 31. The method according to claim 27 wherein the plurality of storage
2 elements are stored on one or more storage systems.

1 32. The method according to claim 31 further comprising:
2 upon determining that a mirrored storage element is to be accessed for the
3 requested file and the mirrored storage element which is to be accessed contains a latest copy
4 of data for the requested file stored on the corresponding master storage element, retrieving
5 the requested file from the mirrored storage element directly.

1 33. The method according to claim 31 further comprising:
2 upon determining that a mirrored storage element is to be accessed for the
3 requested file and the mirrored storage element which is to be accessed does not contain a
4 latest copy of data for the requested file stored on the corresponding master storage element,
5 retrieving the latest copy of data for the requested file stored on the corresponding master
6 storage element from the corresponding master storage element.

1 34. The method according to claim 27 wherein when determining which of
2 the storage elements within the pair of storage elements having the requested file is to be
3 accessed, consideration is given to ensure that all the storage elements within the pair of
4 storage elements having the requested file are substantially accessed in a balanced manner.

1 35. The method according to claim 27 wherein the plurality of storage
2 elements includes a plurality of disk drives.

1 36. A method for optimizing data access between a file server and at least
2 one or more clients, comprising:
3 organizing a plurality of disk drives into pairs for storing a plurality of files,
4 each pair having a master disk drive and one or more mirrored disk drives, and each mirrored
5 disk drive having a copy of data stored on the master drive;
6 maintaining file information on where each of the plurality of files is stored on
7 which pair of disk drives;
8 maintaining access load information on each pair of disk drives, wherein the
9 file information and the access load information are maintained at the file server;
10 upon the file server receiving a request for a requested file from a client,
11 causing the file server to determine which pair of disk drives has the requested file by using
12 the file information, and then causing the file server to determine which of the disk drives
13 within the pair of disk drives having the requested file is to be accessed by using the access
14 load information; and
15 forwarding information relating to the determination as to which of the disk
16 drives within the pair of disk drives having the requested file is to be accessed to client
17 thereby allowing the client to retrieve the requested file.